



US006311205B1

(12) **United States Patent**
Dutcher et al.

(10) **Patent No.: US 6,311,205 B1**
(45) **Date of Patent: *Oct. 30, 2001**

(54) **PERSISTENT USER GROUPS ON SERVERS
MANAGED BY CENTRAL SERVERS**

(75) Inventors: **David Paul Dutcher**, Austin; **William H. Sinclair, Jr.**, Round Rock; **Stanley Alan Smith**, Austin, all of TX (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **09/175,611**

(22) Filed: **Oct. 19, 1998**

(51) Int. Cl.⁷ **G06F 9/00**

(52) U.S. Cl. **709/201; 707/10; 709/224; 709/229; 714/727**

(58) Field of Search **707/10; 709/223, 709/224, 201, 304, 229; 713/202**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,261,102 11/1993 Hoffman 713/200

5,581,749 * 12/1996 Hossain et al. 707/10
5,689,708 * 11/1997 Regnier et al. 709/229
5,748,890 * 5/1998 Goldberg et al. 713/202
5,748,896 * 5/1998 Daly et al. 709/223
5,761,507 * 6/1998 Govett 709/304
6,098,067 * 8/2000 Erickson 707/10
6,209,032 * 3/2001 Dutcher et al. 709/223

* cited by examiner

Primary Examiner—David Wiley

Assistant Examiner—Firmin Backer

(74) *Attorney, Agent, or Firm*—Jeffrey S. LaBaw; Duke W. Yee

(57) **ABSTRACT**

A method for managing user group information in a distributed system comprising a heterogeneous server network. The method enables the establishment of a specified set of user groups on a managed server that are not controlled by a central server. A mechanism is provided to identify the user groups on the managed server, store the groups in a persistent location, and access the user groups during synchronization processing. Synchronization of the persistent user groups at the managed server is done only upon consultation with the persistent account store. The result is a set of user groups on the managed server that can be used to maintain special access and privileges to the persistent user group at the managed server.

3 Claims, 9 Drawing Sheets

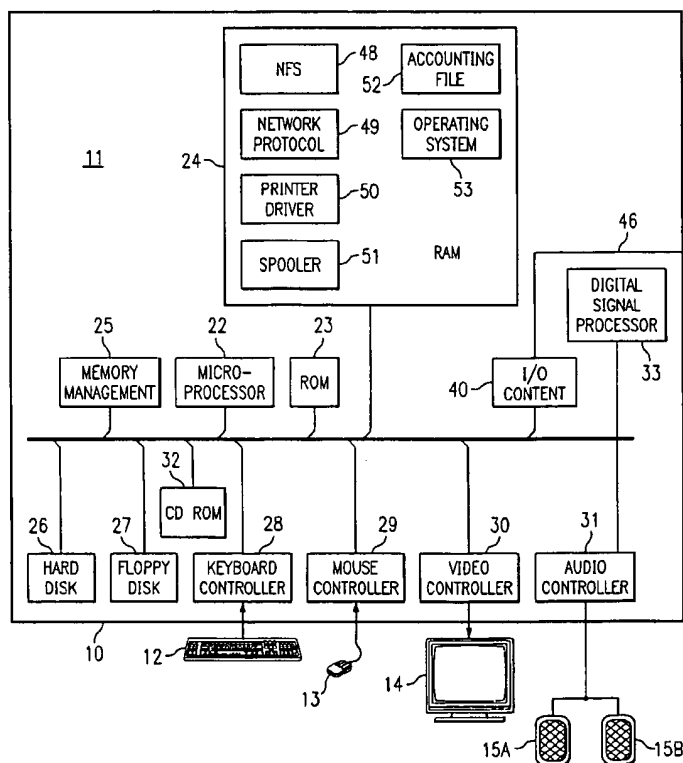
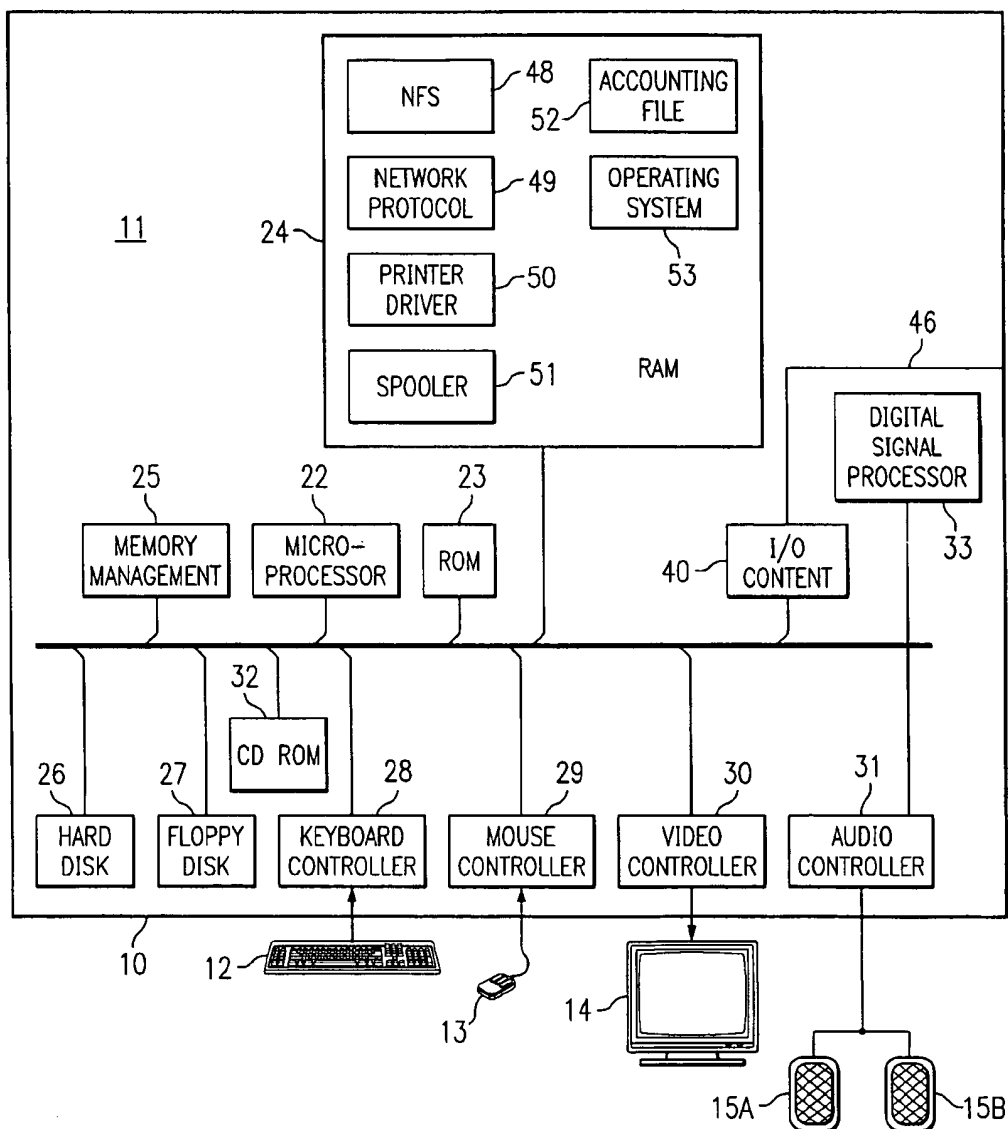
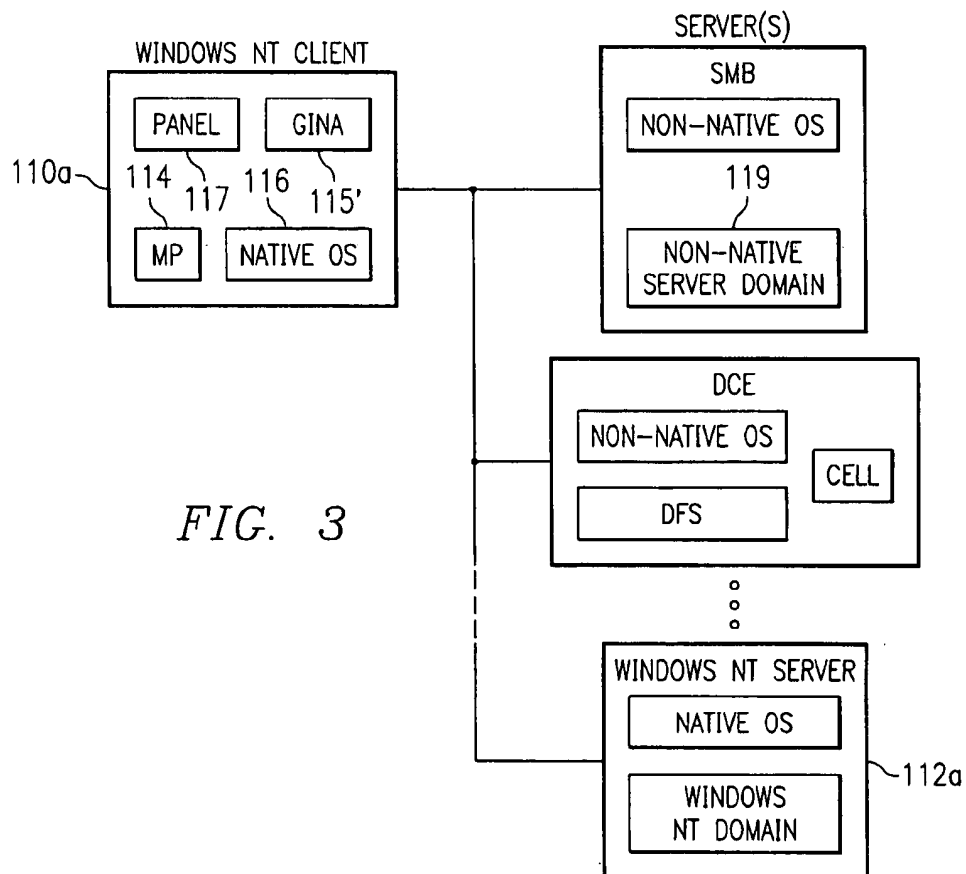
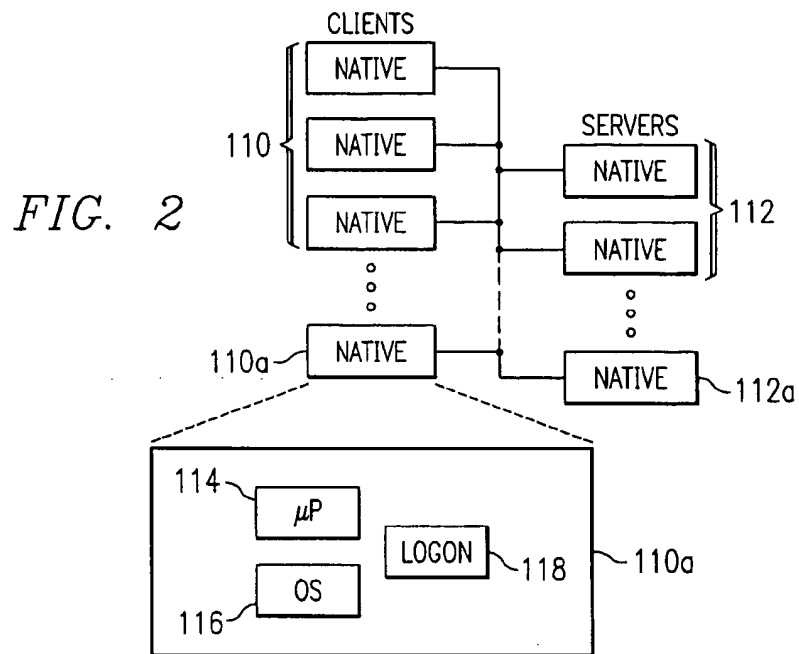


FIG. 1





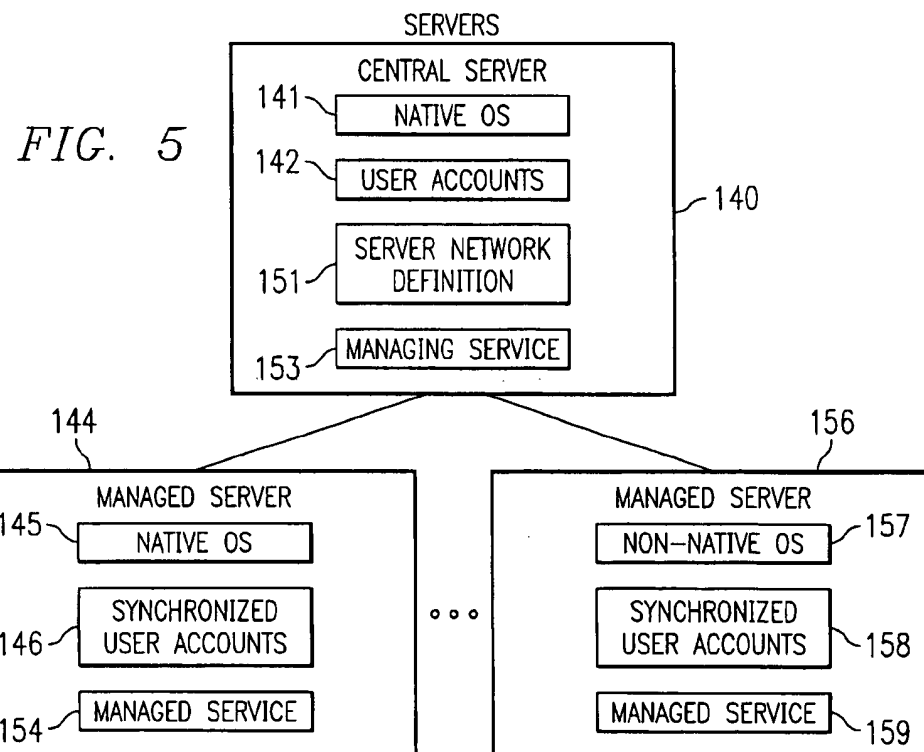
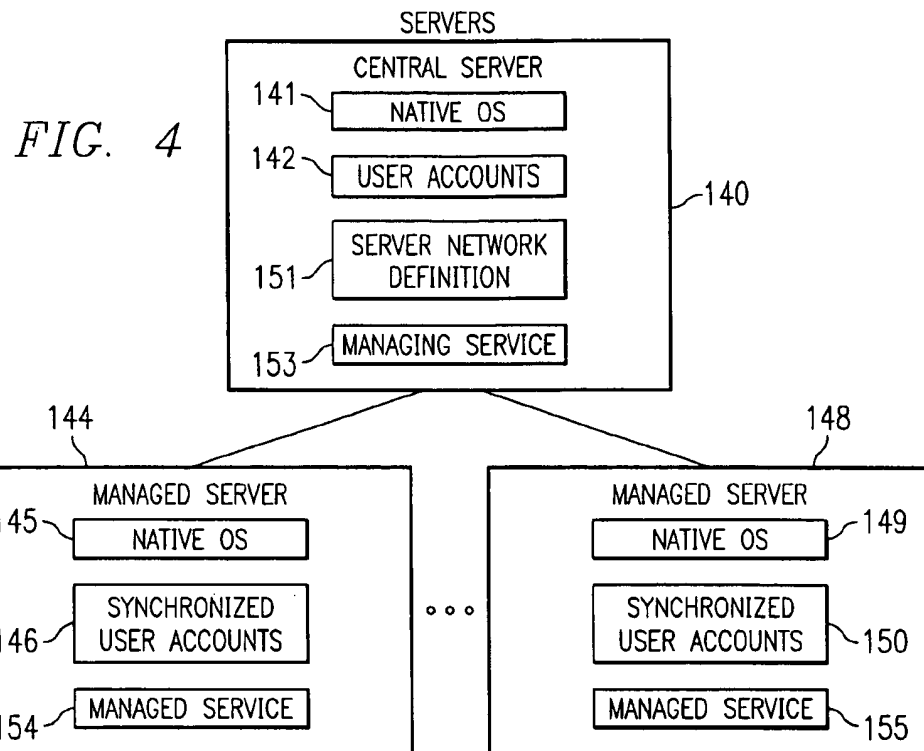


FIG. 6

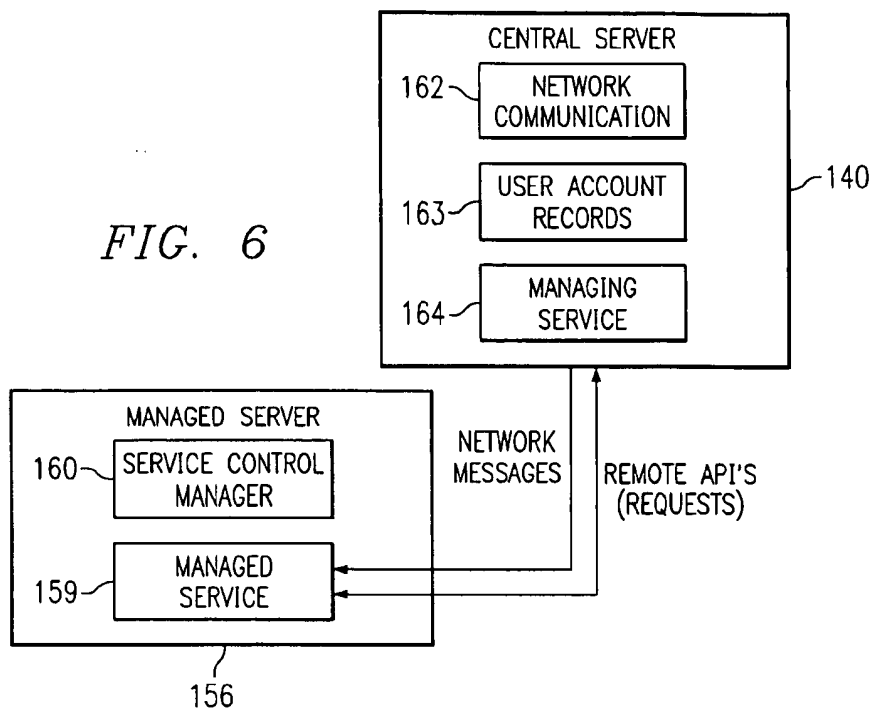


FIG. 7

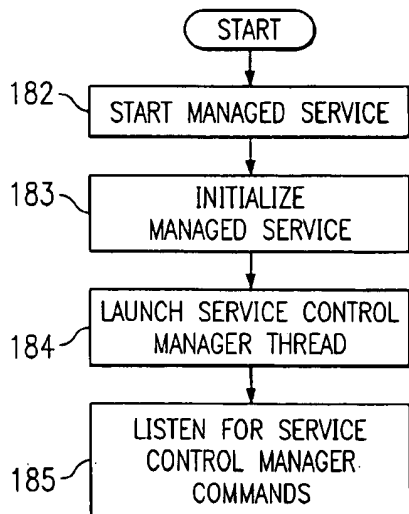


FIG. 8

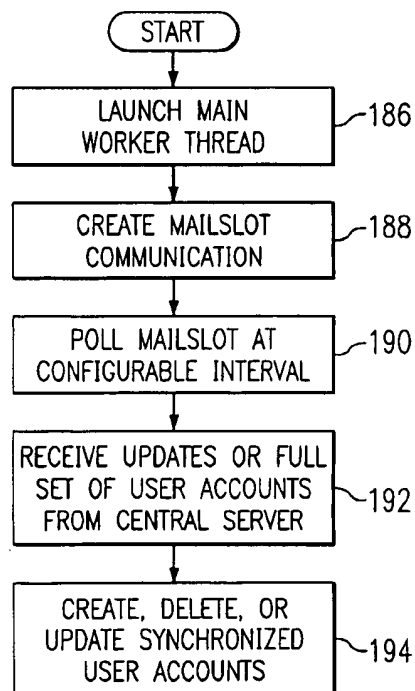


FIG. 9

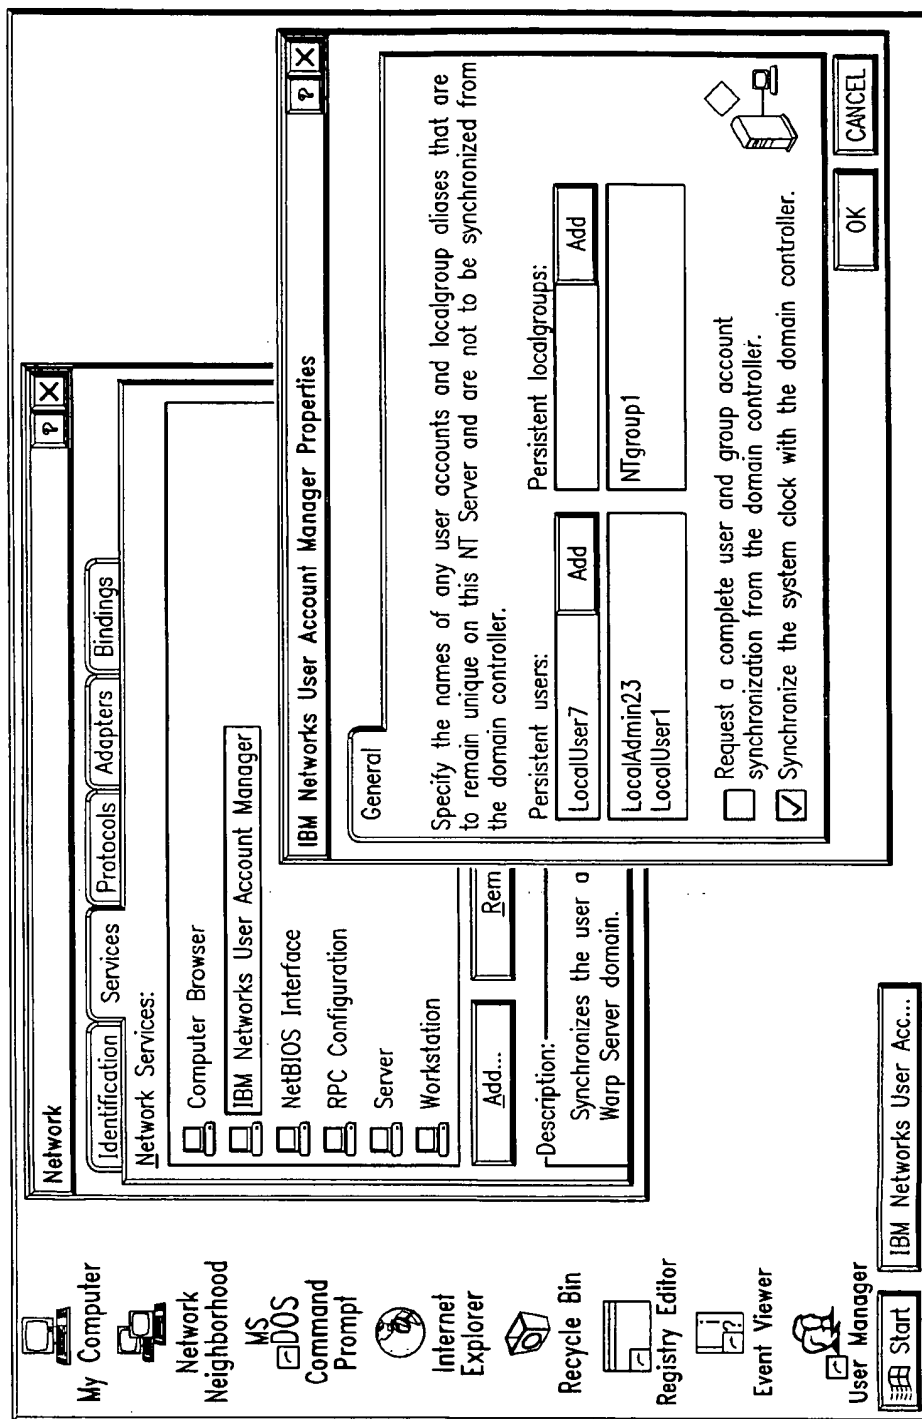


FIG. 10

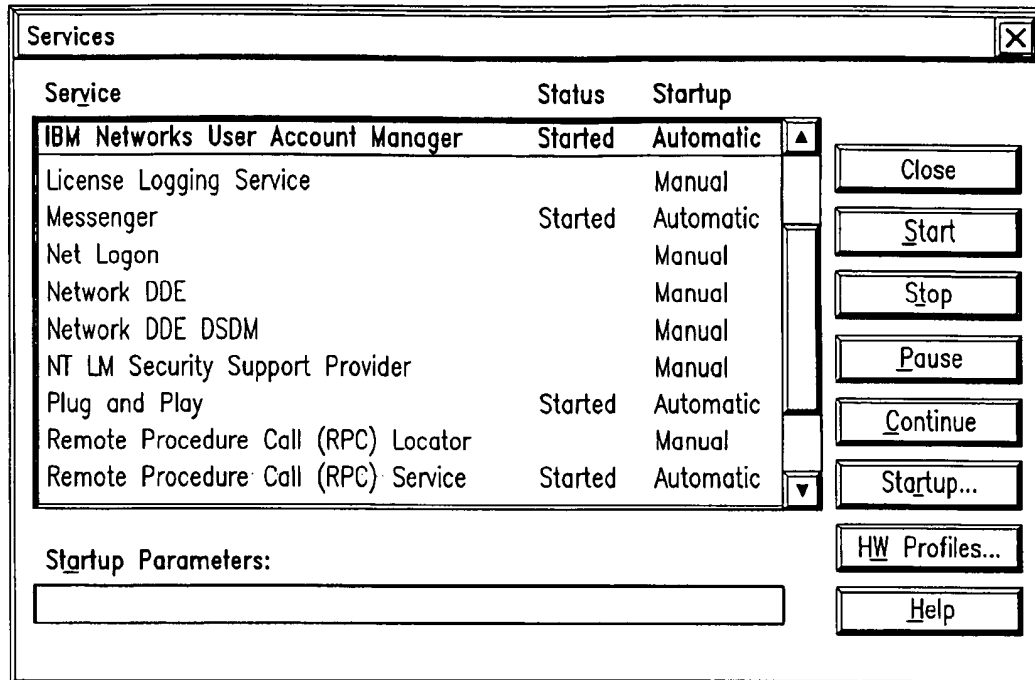


FIG. 11

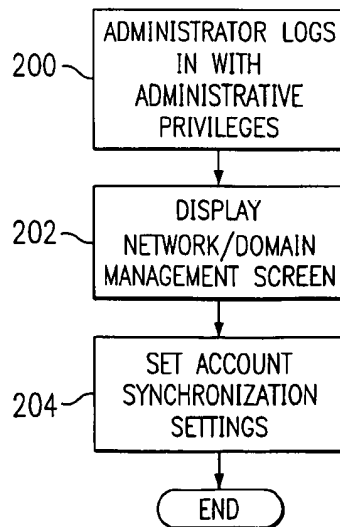


FIG. 12

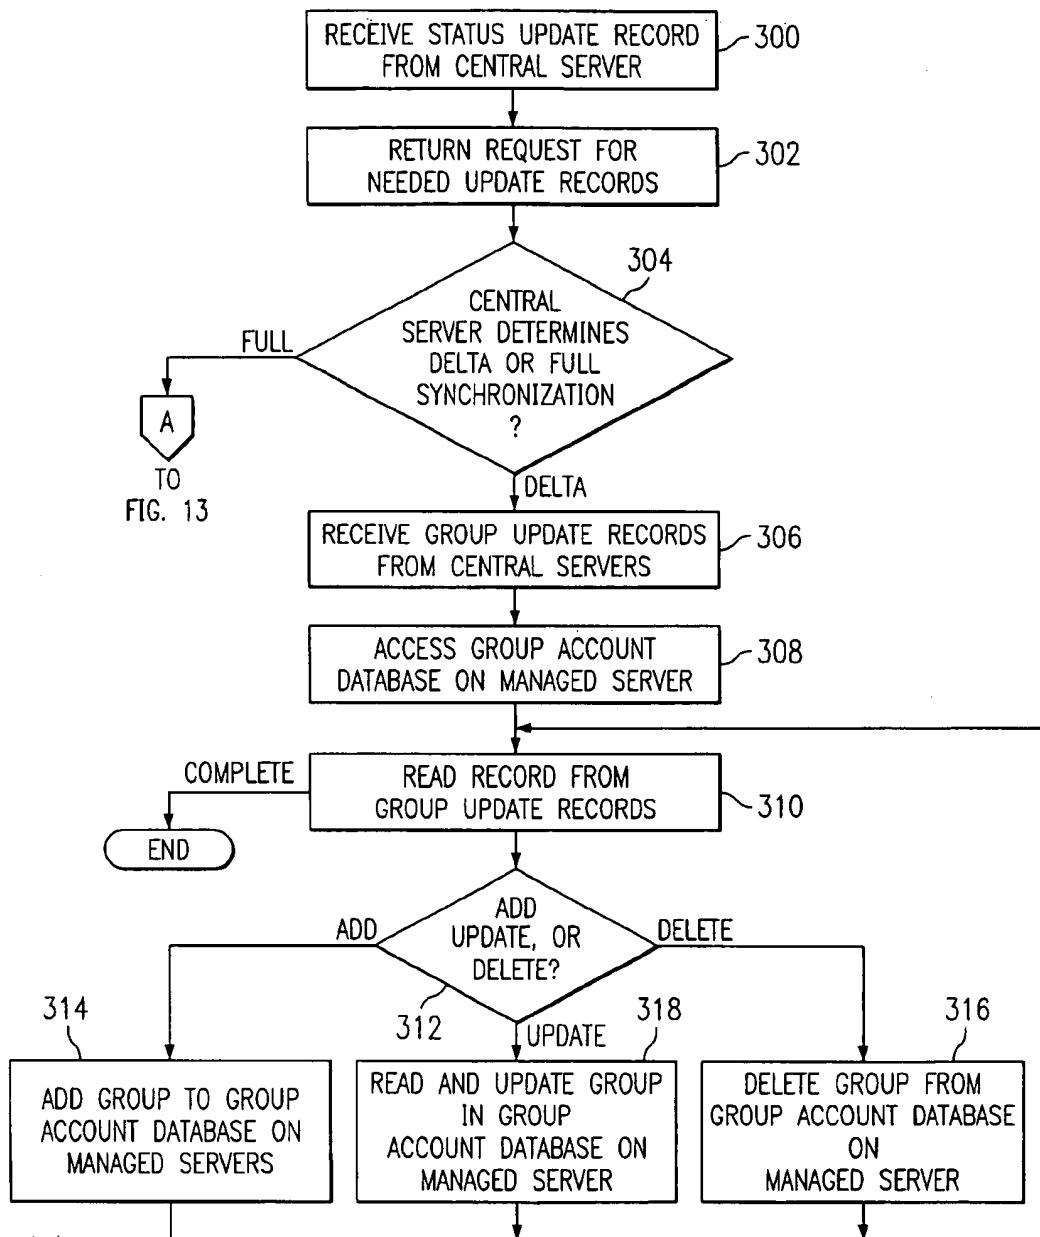


FIG. 13

FROM FIG. 12

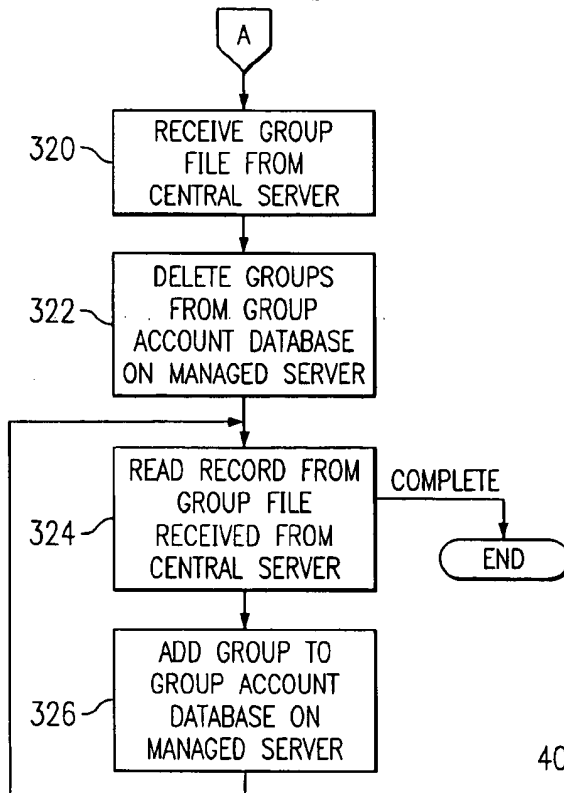


FIG. 14

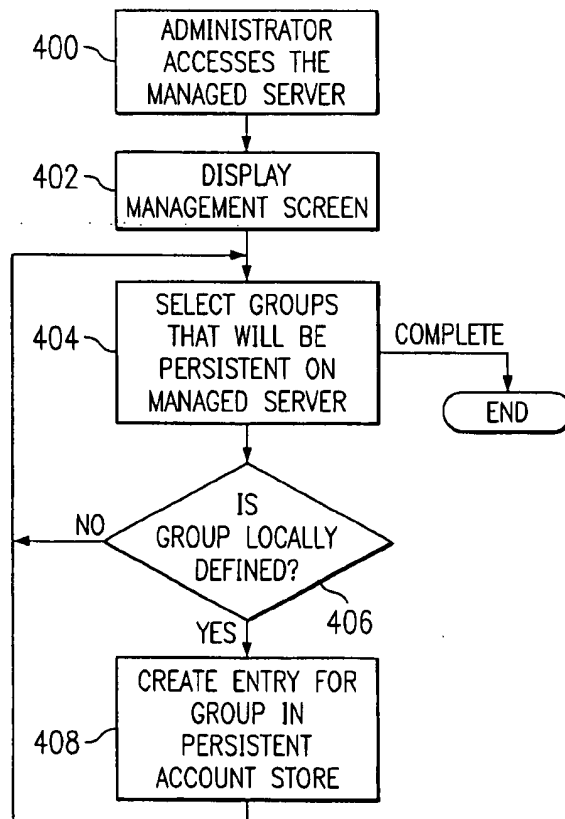
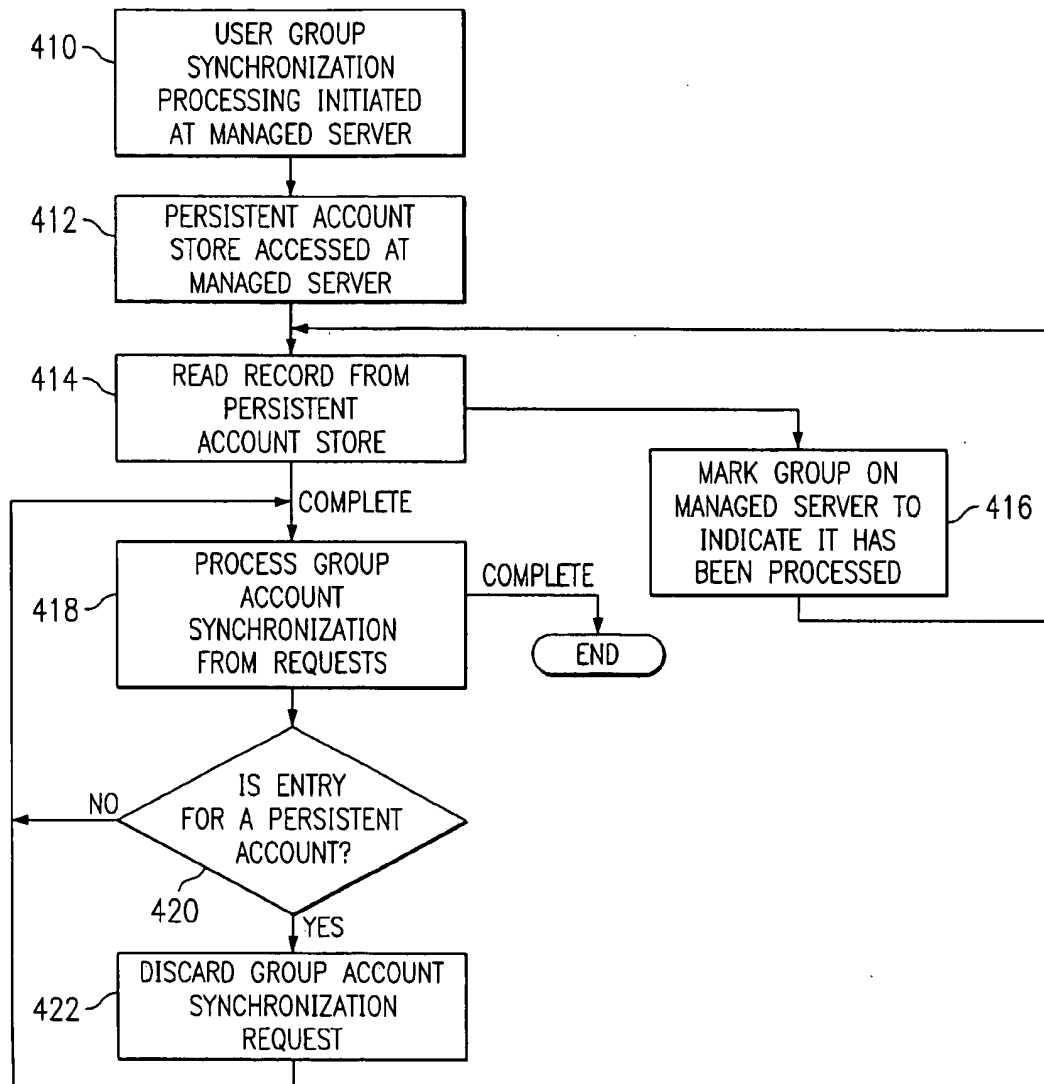


FIG. 15



PERSISTENT USER GROUPS ON SERVERS MANAGED BY CENTRAL SERVERS

BACKGROUND OF THE INVENTION

1. Technical Field

The present invention relates generally to computer networks and more particularly to the segregation of user groups within a server network which consists of a central server and a set of managed servers running native and non-native operating systems which allows for the creation and control of persistent user groups at the managed servers within a heterogeneous server network.

2. Description of the Related Art

The client-server model of computing is a well-known environment. In the model, the user of a computer utilizes a "client" system. The client system runs any of a number of computer operating systems to manage the basic functions that users execute (such as accessing files, executing programs, system administration and the like) as well as to serve as the base against which programs are written. Well-known client operating systems include Microsoft Windows 3.1, Windows for Workgroups, Windows 95, Windows 98, IBM® OS/2® Warp, Apple Macintosh, DOS, many variations of UNIX, and Microsoft Windows NT. The client system serves as the user's workstation, and it may execute programs as well as store some user data.

The server system can also run any of a number of computer operating systems. Well-known server operating systems include Novell Netware, IBM OS/2 Warp Server, IBM AS/400®, Microsoft Windows NT, and many variations of OSF UNIX. The server system is accessed by the client system for specific functions. The functions include, but are not limited to, storage and retrieval of data, storage and execution of applications, and storage of and access to user information.

Server networks are increasingly becoming heterogeneous due to differing problems that can be solved by different servers. User management in these environments requires the creation of different user accounts on the different types of servers. These user accounts eventually have different passwords and possibly different user I.D.'s. A mechanism is needed to allow a single user account definition to be used as the base for any additional user accounts that exist in the network and for a set of user groups to be used as a base for additional user groups in the network. The mechanism needs to go beyond current technology options and allow the accounts and groups on all servers to be continuously updated. While there are advantages to having common user accounts and user groups on the servers, the ability to have a set of user groups on the managed server that are independent of the central server is desirable.

A common term used to refer to a network of related servers is a domain. Within the server domain is a central server acting as the primary domain controller and a plurality of "managed" servers sometimes called secondary servers. Industry standards have been developed (for critical and common functions) to aid in the access from different types of client systems to different types of server systems. The use of these standards on the client and server afford users the opportunity to carry out functions in a consistent manner on a variety of common client and server operating systems. One of the activities that has been standardized is the "authentication" of users. Authentication refers to the process in which a user is validated as being able to complete a log-on and/or access a system. Standard proto-

cols have been defined within the X/Open Server Message Block (SMB) specification and the Open systems Foundation (OSF) Distributed Computing Environment (DCE) specification.

While many products and operating systems have been developed that utilize the standard protocols, not all products have used the standards. When this occurs, either additional work must be done by the other operating system to implement the unique commands used by a vendor, or access to the other new system and/or product is not allowed if the unique commands are not made available to other vendors. When the commands and/or protocol are not made available, that aspect of the system and/or product is sometimes characterized as being "closed". In regards to user management and authentication, the Microsoft Windows NT operating system is becoming an example of a closed server system that is used in many enterprise computer networks.

Server networks are becoming increasingly heterogeneous due to differing application requirements that are solved by different types of servers. Once a server is established, it is desirable to enable access to it without having to manage it completely independent of other servers in the network that are already being utilized. The management of user group access and capability is difficult in a homogeneous server environment where all servers have a common native operating system and, heretofore, it has been virtually impossible to do so in a heterogeneous server environment where servers have native and non-native operating systems. The present invention provides a mechanism to allow establishment of a base set of user groups on a central server in a network that are then used to create and maintain groups across a heterogeneous network of managed servers. The synchronization timing is controlled by the central server and can be done on a real time basis to ensure all servers in the heterogeneous server network are consistent.

While there are many advantages associated with a network of managed servers being controlled by a central server, there are times when it is desirable to keep resources on an additional managed server and secure access to that using a user group based on user accounts or groups known only at that managed server. Such would be the case when multiple common users, say payroll specialists, want to be managed as a group rather than individual users and data and/or functions need to exist on the managed server independent of the central server. A mechanism must exist to designate the user group as a persistent user the managed a specified server that are not updated by the central server in the server network. The invention provides a mechanism to identify the user groups on the managed server and a mechanism to exclude them from synchronization updates sent from the central server. This allows specific user groups on a managed server to have resources in which only they have access or privilege unique to members of the group which remain available independent of central server controlled synchronization.

SUMMARY OF THE INVENTION

It is a general object of this invention to provide segregation of user group account establishment and synchronization on a central server.

It is a more specific object to provide user group synchronization from a central server to a plurality of managed servers in a heterogeneous environment within the server network while preserving specific user groups at any managed server.

3

It is yet another object of the invention to provide specific user groups at managed servers resources exclusive to them. These user groups are independent of central server control.

These and other objects, features and advantages are provided by a method for managing user group information in a distributed system comprising a heterogeneous server network. The method enables the establishment of a specified set of user groups on a managed server that are not controlled by a central server. A mechanism is provided to identify the user groups on the managed server, store the groups in a persistent location, and access the user groups during synchronization processing. Synchronization of the persistent user groups at the managed server is done only upon consultation with the persistent account store. The result is a set of user groups on the managed server that can be used to maintain special access and privileges to the persistent user group at the managed server.q

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and the advantages thereof, reference should be made to the following detailed description taken in connection with the accompanying drawings in which:

FIG. 1 depicts a computer system configured according to the teachings of the present invention;

FIG. 2 is a block diagram of a representative computer network in which the present invention is implemented;

FIG. 3 is a block diagram of the present invention wherein a log-on mechanism is provided in the client running a native operating system to facilitate authentication of a user of the client machine against an account held at a heterogeneous server domain;

FIG. 4 is a block diagram of a homogeneous network of servers;

FIG. 5 is a block diagram of a heterogeneous sever domain having a plurality of managed servers;

FIG. 6 is a block diagram of a central server managing data flow;

FIG. 7 is a flow chart of the managed service start up;

FIG. 8 is a flow chart illustrating managed service operations;

FIG. 9 is a representation of the general properties display screen used to manage services;

FIG. 10 is a screen representation of the managed server definition on a managed server;

FIG. 11 is a flow chart illustrating the setting of central server synchronization values;

FIG. 12 is a flow chart illustrating user group establishment in synchronization in a server domain;

FIG. 13 is a flow chart showing whole group synchronization across a server domain.

FIG. 14 is a flow chart of persistent user group establishment; and

FIG. 15 is a flow chart showing persistent user group utilization.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The invention may be run on a variety of computers or collection of computers under a number of different operating systems. The computer could be, for example, a personal computer, a mini computer, mainframe computer or a computer running in a distributed network of other com-

4

puters. Although the specific choice of computer is limited only by disk and disk storage requirements, computers in the IBM PS/2 (TM) series of computers could be used in the present invention. One operating system which an IBM PS/2 personal computer may run is IBM's OS/2 2.0 (TM).

In the alternative, the computer system might be in the IBM RISC System/6000 (TM) line of computers which run on the AIX (TM) operating system. The various models of the RISC System/6000 is described in many publications of the IBM Corporation. The AIX operation system is described in other publications of the IBM Corporation.

While various computers in the computer networks can be used on the computer network by a client-server setup, FIG. 1 is a common computer interface to the computer network wherein a computer 10 comprising a system unit 11, a keyboard 12, a mouse 13 and a display 14 are depicted in block diagram form. The system unit 11 includes a system bus or plurality of system buses 21 to which various components are coupled and by which communication between the various components is accomplished. The microprocessor 22 is connected to the system bus 21 and is supported by read only memory (ROM) 23 and random access memory (RAM) 24 also connected to system bus 21. A microprocessor in the IBM PC series of computers is one of the Intel family of microprocessors including the 386,486 or Pentium microprocessors. However, other microprocessors including, but not limited to, Motorola's family of microprocessors such as the 68000, 68020 or the 68030 microprocessors and various Reduced Instruction Set Computer (RISC) microprocessors such as the PowerPC chip manufactured by IBM may be used. Other RISC chips made by Hewlett Packard, Sun, Motorola and others may be used in the specific computer.

The ROM 23 contains, among other codes, the Basic Input-Output System (BIOS) which controls basic hardware operations such as the interaction of the processor and the disk drives and the keyboard. The RAM 24 is the main memory into which the operating system and application programs are loaded. The memory management chip 25 is connected to the system bus 21 and controls direct memory access operations including, passing data between the RAM 24 and hard disk drive 26 and floppy disk drive 27. The CD-ROM 32 also coupled to the system bus 21 is used to store a large amount of data, e.g., a multimedia program or presentation.

Also connected to this system bus 21 are various I/O controllers: the keyboard controller 28, the mouse controller 29, the video controller 30, and the audio controller 31. As might be expected, the keyboard controller 28 provides the hardware interface for the keyboard 12, the mouse controller 29 provides the hardware interface for mouse 13, the video controller 30 is the hardware interface for the display 14, and the audio controller 31 is the hardware interface for the speakers 15. An I/O controller 40 such as a Token Ring Adapter enables communication over a network 46 to other similarly configured data processing systems.

As described above, one of the preferred implementations of the invention is as sets of instructions 48-52 resident in the random access memory 24 of one or more computer systems configured generally as described above. Until required by the computer system, the set of instructions may be stored in another computer readable memory, for example, in the hard disk drive 26, or in a removable memory such as an optical disk for eventual use in the CD-ROM 32 or in a floppy disk for eventual use in the floppy disk drive 27. Further, the set of instructions can be

5

stored in the memory of another computer and transmitted in a transmission medium over a local area network or a wide area network such as the Internet when desired by the user. One skilled in the art would appreciate that the physical storage of the sets of instructions physically changes the medium upon which it is stored electrically, magnetically, or chemically so that the medium carries computer readable information. While it is convenient to describe the invention in terms of instructions, symbols, characters, or the like, the reader should remember that all of these and similar terms should be associated with the appropriate physical elements.

Further, the invention is often described in terms that could be associated with a human operator. While the operations performed may be in response to user input, no action by a human operator is desirable in any of the operations described herein which form part of the present invention;

FIG. 2 illustrates a computer network having one or more "client" machines 110 and one or more "servers" 112. A typical client machine 110a is a personal computer or workstation running an Intel processor 114 and the Microsoft Windows NT 4.0 operating system 116. For convenience herein, a machine configured in this manner is sometimes referred to as a "Windows NT client". Any other type of hardware platform that runs Windows NT operating system may be used as the client. According to the present invention, the client also includes an application 118, which provides certain additional functionality to achieve the objects of the present invention. Each client has basic networking hardware to establish a connection out to a server. Thus, for example, a client may have a TCP/IP or NETBIOS connection to the network running over a token ring or Ethernet adapter.

Typically, a server in the computer network is another personal computer or workstation platform that is Intel, Power PC® or RISC® based, and includes an operating system such as Windows NT 4.0, IBM® OS/2® Warp Server, AIX® or the like. At least one server 112a in the computer network is the central server and executes the base operating system or subsystem which is termed "native". This "native" system could be an IBM OS/2 Warp Server, which is sometimes referred to as a "Warp Server". A server 112 is said to be "native" if it is running the same operating system as the server 112a. A "non-native" server is thus a server platform (e.g., a personal computer) running an operating system or subsystem that is different than the operating system running on the server system 112a. Given an IBM OS/2 Warp Server as 112a, examples of such "non-native" servers include, without limitation, Microsoft Windows NT Server, Novell Netware Server, other types of server Message Block (SMB) servers, as well as operating systems that run Open Systems Foundation (OSF) Distributed Computing Environment (DCE) software. An example of the latter is a DCE Cell running Distributed File System (DFS).

In the prior art, a mechanism exists to allow a user at client system 110a to authenticate to a server domain using a user account held at that domain. As seen in FIG. 3, the module GINA 115' ("graphical identification and authorization") is registered on the example Windows NT client. This enables the Windows NT client user to be authenticated against an account held at a native or non-native server domain 119. As used herein, a "non-native server domain" refers to a database of user account information retained at a given server running an operating system that is different than the operating system running at the client system. The term "heterogeneous client-server

6

network" is commonly used to describe an environment in which the client operating system and server operating system are different. This type of environment is common in the client-server model. In contrast, the term "homogeneous client-server network" is commonly used to describe an environment in which the client operating system and server operating system are the same.

A non-native server domain is typically supported on a non-native server. Thus, where the user seeks authentication from a Windows NT client, a non-native server domain includes, without limitation, any Server Message Block (SMB) server domain (e.g., IBM Warp Server 4.0), a DCE Cell in which Distributed File System (DFS) is implemented, or other known domains such as UNIX domains. This is illustrated in FIG. 3. Of course, the computer network can also include a Windows NT server domain 112a if authentication is sought from a native server domain.

While "heterogeneous client-server networks" and "homogeneous client-server network" are part of the environment for this invention, they are not specific to it. This invention is specific to the network of servers. In particular, it refers to a "server network" consisting of a central server and a set of managed servers. The managed servers can be native to the central server or non-native to the central server.

In the described embodiment, the central server is an OS/2 Warp Server managing a heterogeneous mix of OS/2 servers and Windows NT servers. One of ordinary skill in the art should appreciate that teachings of this invention are also applicable where the server is running other types, native or non-native servers within the server domain. The present invention operates in a network where a single user account or user group from a central server is used as the basis to create the corresponding user account or user group on different type of managed server. After establishment of the user account or group, the central server where the initial account was established will ensure all account updates are replicated to the managed servers on a real time basis.

FIG. 4 is a block diagram representing a high level operation of the present domain server network having a central server 140 with a native operating system 141, user accounts 142, server network definition 151 and managing service 153. Central server 140 controls homogeneous managed servers 144 and 148. Within managed server 144 is a native operating system 145, synchronized user accounts 146 and managed service 154. Likewise managed server 148 has a native operating system 148, synchronized user accounts 150 and managed service 155. The central server 140 can manage a multiplicity of servers such as servers 144 and 148 having a native OS operating system and synchronized user accounts and managed service.

The server network is defined in the central server by server network definition 151. The server network definition 151 provides the network definition required to link the servers together. It consists of a table of "addresses" that are used to direct and receive communications between specific servers in a network. Whether a broadcast mechanism like NETBIOS or a more directed mechanism like TCP/IP is used for communication, the appropriate network addresses would be included in the table. One skilled in the art would recognize that portions or all of the server network definition could be stored at the managed server(s) and/or the central server.

Similarly the managing service 153 is used to control a managed service 154 within managed server 144 and 148.

The managing service 153 and managed services 154 and 155 provide the mechanism to synchronize user accounts.

Managing service 153 and managed service(s) 154 and 155 are the core components used to manage the user accounts across the servers. In a "homogeneous server network" where the servers utilize the same native operating system, the services are implemented through usage of common support that exists on each server. In addition, user accounts 142 and synchronized user accounts 146 are commonly created, updated, and deleted while containing the same information in each.

Prior art exists for user account synchronization in the "homogeneous server network". User account management commands are sent between the central server 140 and the managed server(s) 144 and 148. Since the same commands (application programming interfaces) exist on both servers in a "homogeneous server network", the services are there to support the remote command execution between servers. An example of this is an IBM OS/2 Warp Server created as a Primary Domain Controller which would be a central server and an IBM OS/2 Warp Server created as an additional server which would be a managed server.

Shown in FIG. 5, central server 140 having a required operating system 141, user accounts defined in 142, a server network definition 151 and a managing service 153. The central server 140 now manages a managed server 144 having a native operating system 145, synchronized user accounts 146, and a managed service 154, as well as heterogeneous server 156 having a non-native operating system 157, synchronized user accounts 158 and a managed service 159. In the described embodiment, the native operating system is an IBM OS/2 Warp Server and the non-native operating system is a Microsoft Windows NT server.

Similar to the "homogeneous server network", managing service 153 and managed service(s) 154 and 159 are the core components used to manage the user accounts across the servers in the "heterogeneous server network". In a "heterogeneous server network" where the servers utilize native OS such as 145 and non-native OS such as 157, common support does not exist on each server. With no common support guaranteed between the central server and the managed server(s), a new and unique mechanism is required to enable the establishment and synchronization of user accounts between the unlike servers.

Managing service 153 and managed service 159 represent the situation where the servers are different with one using a native OS and one using a non-native OS. Managing service 153 is responsible for determining what each managed server requires and delivering it in a format that can be used at the managed server 156. This involves the tracking of user account changes, packaging of the changes in a known format, and sending them to the correct managed server when required. In the described embodiment, the central server 140 is an IBM OS/2 Warp Server and the user account changes are packaged in the same format as used to synchronize to a managed server 144 that is also an IBM OS/2 Warp Server.

Managed service 159 is responsible for implementation of support to maintain communication with central server 140. The service receives the account changes, determines what has been changed with the user account, and implements the correct set of application programming interface calls to effect the changes for the synchronized user accounts 158 on the managed server 156. The mechanism to do this will be different for each non-native OS and is detailed in later figures.

In operation, central server 140 sends information called mailslot broadcasts to the network from a network communication mechanism 162, as shown in FIG. 6, to the managed server 156, while also responding to requests received from managed servers in the network. Through the network, network messages are used to provide requests and responses to the managed service 159 within the managed server 156. A user account subsystem containing a database of records is maintained in user account records 163. Also resident in the central server 140 is the managing service 164 that was described in detail as mechanism 153 in FIG. 5. Central server 140 and managed server 156 make specific functional requests through application programming interfaces (API) which specify requests and responses as described in FIG. 5.

To enable the receipt and processing of requests sent from the central server, a mechanism is required to start the managed service on the managed server. Once started, the managed service continuously awaits status updates and requests. Shown in FIG. 7, a service control manager is required on the managed server. Either when the managed server is started or through a request for a user with administrative capabilities, the managed service is started in step 182. Following the starting of the service, settings that control the operation of the service are initialized in step 183. These settings are described in FIG. 12. To allow the service to respond to commands issued by the administrator a mechanism must be established. In step 184, a thread in the operating system is launched to allow the service to recover commands as shown in step 185. Commands include the changing of settings, stopping the service, and starting the service.

At the same time that the service control manager thread is launched in step 184 of FIG. 7, the main operational support must also be started on the managed server. FIG. 8 shows the steps to establish operational support. A thread or equivalent mechanism is started on the managed server as in step 186. A mailslot or similar communication mechanism is then established at the server in step 188 to enable communication with the central server. To monitor status from the central server, the mailslot or similar communication mechanism must be checked on a regular interval. Step 190 includes the process of checking on a regular basis. For each server in the heterogeneous server network, the communication support must be implemented as appropriate for the non-native OS. Steps 192 and 194 introduce the receipt and processing of user account requests from the central server. The main worker thread shown in step 186 is used to execute those steps which are detailed in FIG. 12 through FIG. 15, and are implemented for each heterogeneous network server.

The establishment of the managed service at the network server can be through local administration as shown in the screen printout of the managed service screen shot shown in FIG. 9 and FIG. 10. To define the service which must be active to enable the ongoing synchronization of user accounts, the service is installed on the server that is to be managed. The method and display of the service may differ by the type of heterogeneous server in the network, and FIG. 9 shows the service in the described embodiment of a Microsoft Windows NT server as a managed server. Communication between a central server and the managed servers is required. In FIG. 9 the managed service which is called "IBM Networks User Account Manager", is established as part of the Network portion of the server. Additional settings can be established through properties associated with the managed service. Synchronization of the clock on the managed server with the central server clock is one of the settings that is made.

9

FIG. 10 shows additional controls used for the managed service on the managed server. For synchronization of user accounts to be in effect, the managed service must be active. In FIG. 10, the screen shot of the screen used by the administrator to control the service is shown. The method to start can either be manual or automatic. The status of the service should also be able to be retrieved to allow the administrator to know whether synchronization is active. The status column in FIG. 10 is an example of a status display.

To establish and control the managing service on the central server, the server must be accessed with administrative privilege as shown in step 200 of FIG. 11. The managing service needs to exist on the server and in the case of the described embodiment of the central server being in IBM OS/2 Warp Server, it is the LAN manager service used for homogeneous servers. Step 202 shows the access to the settings that control operation of the managing service. The actual settings are done in step 204. Settings that are needed include the interval that status updates (described as pulses) are exchanged from the central server to the managed server(s), and how much difference is allowed between the servers. The combination of these settings determine how frequently updates occur to synchronize the user accounts between servers.

The present invention addresses the problem of managing user group access and capability in a server network that contains a heterogeneous set of native and non-native servers. Server networks are becoming increasingly heterogeneous due to differing application requirements that are solved by different types of servers.

For user group establishment and synchronization, the groups must first be created on the central servers. This invention is not directed to the initial creation or updating of groups on the central server. Existing means that are part of the native operating system or subsystem are used for the creation and update.

To establish user groups and provide for synchronization, the managed server receives a status update record from the central server as in step 300 in FIG. 12. A return request for needed update records is then sent from the managed server to the central server as shown in step 302. The central server determines whether a delta update should be done since the last update or if a full synchronization needs to be made in step 304. The determination should be based on the number of records that will be required to update the managed server. If the number of updates is large, it should be determined that a full synchronization will be done. If only an incremental or delta update needs to be made, the managed server receives user group update records from the central server as per step 306. Since the user group format and method to define and access user groups differs for each type of server, the managed service is responsible for determining and issuing the API calls to access the user groups and read the individual records. Access to the user group records on the managed server is then made in step 308. This is done by the managed server when it determines account updates are required. The record is read from group update records in step 310 by the managed service using API's that exist on the heterogeneous managed server.

Once the record is read, the managed service makes the determination in step 312 to add, update, or delete the user group account. If the request is to add, the service is responsible for using the API call appropriate on the server to add the user group to the user account database on the managed server as shown in step 314. If the request is to

10

update, step 318 is executed and the service issues the appropriate API call on the managed server causing the user group to be read and updated on the managed server. If the request is to delete, step 316 is executed and the service issues the appropriate API call or the managed server to cause the user group to be deleted from the user account database on the managed server. This series of operations continues on the managed server until the read records from group update records 310 is complete which occurs when all records have been processed.

The managed service is responsible for receiving the user group records file in step 320 of FIG. 13 that contains all user groups from the central server. To ensure all group definitions on the managed server are consistent with the group definitions on the central server, the managed service running on the managed server issues the API calls appropriate for the managed server to delete all user groups from the group account database on the managed server as shown in step 322. Following deletion of the groups, each of the records received from the central server are read in step 324. The managed service is then responsible for issuing the API call appropriate for the managed server to create a group based on the information received from the central server as shown in step 326. The group creation includes associating the user accounts defined as part of the group with the group definition. Each record is read and processing is complete when the last record has been processed.

While central control of a server network is often desirable, there are situations when a local administrator has some secure data they do not want to disturb or the administrator may need a user group that is unique to the managed server for unique administrative control. For these and other situations, there is a need to have specific user groups kept constant at a managed server. The present invention creates a mechanism in which an individual with the proper privileges at the managed server can designate one or more user groups that will be managed independent of the central server. This enables all general user groups on the managed server to be controlled by the central server while preserving specific user groups at the managed server. Special user actions or access are maintained at the managed server and associated with the persistent user groups on the server.

This is accomplished by referring to FIG. 14 wherein an individual with administrative privilege accesses the managed server in step 400. In step 402 a management screen is displayed to enable the administrator to select the persistent user groups. The administrator in step 404 selects the set of user accounts that will be persistent on the managed server. FIG. 9 shows an example of the management screen used in the described embodiment. The user groups are selected in the box titled "Persistent Groups." An entry is created on the managed server for the group in a persistent account store as shown in step 408. If the user group is not locally defined on the managed server then it is not made persistent and it continues to be maintained by the central server. After all user groups that have been selected to be persistent are processed, the actions are completed.

The persistent user groups must be accessed during synchronization initiated by the central server. These accounts must be excluded from the changes that are potentially requested for a user group. Step 410 in FIG. 15 reflects the initiation of the synchronization processing at the managed server. Step 300 of FIG. 12 is an example of the initiation of synchronization processing.

To determine the user groups that are persistent on the managed server, the persistent account store must be

11

accessed on the managed server as in step 412. Each entry or record from the persistent account store on the managed server must be read. Step 414 shows the reading of a record from the persistent account store. Prior to delta update processing as shown in step 306 of FIG. 12 or the full synchronization processing as shown in step 320 of FIG. 13, the user groups on the managed server must be masked. By masking the user group records, the managed service can exclude them during synchronization processing which will retain the existing user group information.

During the user group synchronization, each user account synchronization request is processed as shown in step 418. When the user group request received from the central server matches a persistent user group on the managed server as shown in step 420, the user group synchronization request is discarded. By discarding the changes for the user groups, they remain persistent at the managed server and are unchanged by central server actions.

While the invention has been shown and described with reference to particular embodiments thereof, it will be understood by those skilled in the art that the invention can be practiced, with modification, in other environments. For example, although the invention described above can be conveniently implemented in a general purpose computer selectively reconfigured or activated by software, those skilled in the art would recognize that the invention could be carried out in hardware, in firmware or in any combination of software, firmware or hardware including a special purpose apparatus specifically designed to perform the described invention. Though two managed servers were shown and described a multiplicity of managed servers may be controlled by a central server. Therefore, changes in form and detail may be made therein without departing from the spirit and scope of the invention as set forth in the accompanying claims.

What is claimed is:

1. In a client server computer network comprised of a central server, and a plurality of managed servers, and in which network full synchronization of managed server records is controlled by said central server, a method for synchronization of user records on managed servers with user records on said central server, including the steps of:

 sending a status update pulse from the central server to managed servers on the network;

12

 sending, in response to said pulse update record, a return request from a managed server to the central server for update records;

 sending, in response to said return request, from said central server to said managed server, an update record request or a request for full synchronization;

 defining at the managed server a set of user records locally on the managed server to be persistent accounts on this managed server;

 excluding at the managed server from potential changes on the managed server, as a result of a central server full synchronization request, those user records that have been defined as persistent accounts.

2. The computer as recited in claim 1, wherein the native operating system is IBM OS/2 and the non-native operating system is Windows NT.

3. In a client server network comprised of a central server, a plurality of managed servers at least some of which managed servers are non-native with respect to said central server, and in which network user group records are established and are updated on said central server, a method for synchronization of user group records on managed servers in said network with user group records on said central server, including the steps of:

 sending a status update pulse from the central server to managed servers on the network;

 sending, in response to said status update pulse, a return request from a non-native managed server to the central server for update records;

 sending, in response to said return request, from said central server to said non-native managed server, an update records request or a request for full synchronization;

 defining at the non-native managed server a set of user records locally on the non-native managed server to be persistent accounts on this non-native managed server;

 excluding at the non-native managed server from potential changes on the non-native managed server, as a result of a central server full synchronization request, those user records that have been defined as persistent accounts.

* * * * *